

UTILITY PATENT APPLICATION TRANSMITTAL

Submit an original and a duplicate for fee processing
(Only for new nonprovisional applications under 37 CFR 1.53(b))

ADDRESS TO:

Assistant Commissioner for Patents
Box Patent Application
Washington, D.C. 20231

Attorney Docket No. 99,360

First Named Inventor Janakiraman
Senthilnathan

Express Mail No. EL007910479US

Total Pages

APPLICATION ELEMENTS

1. ☒ Transmittal Form in duplicate with Fee
2. ☒ Specification (including claims and abstract) [Total Pages 24]
3. ☒ Drawings [Total Sheets 8]
4. ☐ Oath or Declaration [Total Pages]
 - a. ☐ Newly executed
 - b. ☐ Copy from prior application

[Note Boxes 5 and 17 below]

 - i. ☐ Deletion of Inventor(s) Signed statement attached deleting inventor(s) named in the prior application
5. ☐ Incorporation by Reference: The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. ☐ Microfiche Computer Program
7. ☐ Nucleotide and/or Amino Acid Sequence Submission
 - a. ☐ Computer Readable Copy
 - b. ☐ Paper Copy
 - c. ☐ Statement verifying above copies

ACCOMPANYING APPLICATION PARTS

8. ☐ Assignment Papers
9. ☐ Power of Attorney
10. ☐ English Translation Document (if applicable)
11. ☐ Information Disclosure Statement (IDS)
 - ☐ PTO-1449 Form
 - ☐ Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Return Receipt Postcard (Should be specifically itemized)
14. ☐ Small Entity Statement(s)
 - ☐ Enclosed
 - ☐ Statement filed in prior application; status still proper and desired
15. ☐ Certified Copy of Priority Document(s)
16. ☒ Other: Certificate of Express Mail

17. If a **CONTINUING APPLICATION**, check appropriate box and supply the requisite information:
☐ Continuation ☐ Divisional ☐ Continuation-in-part of prior application Serial No. .

APPLICATION FEES

BASIC FEE				\$760.00
CLAIMS	NUMBER FILED	NUMBER EXTRA	RATE	
Total Claims	19	-20=	x \$18.00	\$
Independent Claims	3	- 3=	x \$78.00	\$
<input type="checkbox"/> Multiple Dependent Claims(s) if applicable			+\$270.00	\$
Total of above calculations =				\$760.00
Reduction by 50% for filing by small entity =				\$()
<input type="checkbox"/> Assignment fee if applicable			+\$40.00	\$
TOTAL =				\$760.00

UTILITY PATENT APPLICATION TRANSMITTAL

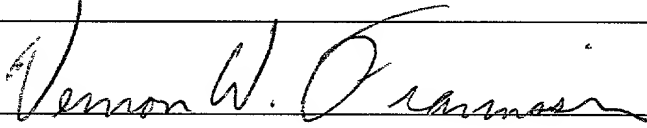
Attorney Docket No. 99,360

18. ☐ Please charge my Deposit Account No. 13-2490 in the amount of \$.
19. ☒ A check in the amount of \$760.00 is enclosed.
20. The Commissioner is hereby authorized to credit overpayments or charge any additional fees of the following types to Deposit Account No. 13-2490:
- a. ☒ Fees required under 37 CFR 1.16.
- b. ☒ Fees required under 37 CFR 1.17.
- c. ☒ Fees required under 37 CFR 1.18.
21. ☒ The Commissioner is hereby generally authorized under 37 CFR 1.136(a)(3) to treat any future reply in this or any related application filed pursuant to 37 CFR 1.53 requiring an extension of time as incorporating a request therefor, and the Commissioner is hereby specifically authorized to charge Deposit Account No. 13-2490 for any fee that may be due in connection with such a request for an extension of time.

22. CORRESPONDENCE ADDRESS

Name	McDonnell Boehnen Hulbert & Berghoff
Address	32 nd Floor, 300 South Wacker Drive
City, State, Zip	Chicago, Illinois 60606

23. SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT REQUIRED

Name	Vernon W. Francissen Registration No. 41,762
Signature	
Date	October 29, 1999

UTILITY (Rev. 11/18/97)

CERTIFICATE OF MAILING BY "EXPRESS MAIL"
(NEW PATENT APPLICATION)

Express Mail No. EL007910479US

Deposited October 29, 1999

I hereby certify that the attached correspondence, identified below, is being deposited with the United States Postal Service as "Express Mail Post Office to Addressee" under 37 CFR § 1.10 on the date indicated above and is addressed to the Asst. Commissioner for Patents, Box New Application, Washington, DC 20231.

By: Frank Clarke
(person actually depositing)

Patent Application of: Janakiraman Senthilnathan, Rohit Verma, Amit Tiwari, Ashwin Hegde

Title: A Method and Apparatus for Selection of an Endpoint Device in a Point to Point Link

1. Utility Patent Application Transmittal in duplicate;
2. Specification including claims and Abstract (24 sheets);
3. Drawings (8 sheets);
4. Filing Fee Check in the amount of \$760.00; and
5. Return postcard.

Attorney Docket No.: 99,360

APPLICATION FOR A UNITED STATES PATENT
UNITED STATES PATENT AND TRADEMARK OFFICE
(CASE No. 99,360)

**Title: A METHOD AND APPARATUS FOR SELECTION OF AN ENDPOINT
DEVICE IN A POINT TO POINT LINK.**

5 Inventors: Janakiraman Senthilnathan, a citizen of India, and a resident of Mount Prospect,
 Illinois.

 Rohit Verma, a citizen of India, and a resident of Schaumburg, Illinois.

10 Amit Tiwari, a citizen of India, and a resident of Schaumburg, Illinois.

 Ashwin Hegde, a citizen of India, and a resident of Arlington Heights, Illinois.

15

Assignee: 3Com Corporation
 5400 Bayfront Plaza
 Santa Clara, CA 95052

20

FIELD OF INVENTION

The present invention relates to data communications. More specifically, it relates to the transmission of packets in a point to point communication link.

5

BACKGROUND OF THE INVENTION

Connection oriented point-to-point communication links, such as a Layer 2 Tunneling Protocol (L2TP) tunnel, are an increasingly common feature of network infrastructures. Tunnels are prearranged connections established by agreement between internet service providers (ISPs). See Request for Comment (RFC) 2661 and *Layer Two Tunnelling Protocol (L2TP)*, A. Valencia, et al., draft-ietf-pppext-l2tp-16.txt, June 1999, herein incorporated by reference, available from the Internet Engineering Task Force (IETF) at www.ietf.org for more information. FIG. 1 shows an architecture 10 that illustrates two L2TP tunnels 56 and 66 established through a public IP network 70. Each L2TP tunnel is a prearranged point to point link between remote client 20 and server 80.

An L2TP tunnel provides a conduit for communications between a remote client 20 and a server 80. Typically, a single tunnel slot provides the communication link between a client and server. However, it is increasingly common for there to be multiple tunnels providing the client-server communication link. For example, Multi-Link Point-to-Point (MLPPP) connections aggregate the bandwidth of multiple tunnel connections to provide a single higher bandwidth communication link. Also, in wireless mobile applications, a second tunnel link may be established from a tunnel initiator in the cell area that the client is entering while a first tunnel

link still exists from the tunnel initiator in the cell area that the client is leaving. It is advantageous in such multiple link connections to have each link terminate on the same tunnel endpoint. However, there is no conventional way to deterministically select the endpoint for a given L2TP tunnel.

5 When remote client 20 establishes a dial-up connection 22 or 24 with a tunnel initiator (TI) 30 or 40, then the TI typically recognizes client 20 as a tunnel client by means of an authentication protocol, such as RADIUS, see Request For Comment (RFC) 2138, herein incorporated by reference, or other means for identifying the client. Typically, each TI has a table that indicates the endpoint for the tunnel connection for client 20. The table in each TI
10 typically includes a list of tunnel endpoints (TEs), such as 50 and 60 for each remote client and each TI selects an endpoint from the list independent of the selection made by another TI. Similarly, TI 40 will have a table that indicates a list of endpoints for client 20. When client 20 connects to TI 30 or 40, then each TI will independently select a TE device. As a result, there is a high likelihood that the two tunnel connections 56 and 66 will terminate on different tunnel
15 endpoint devices, as is shown in FIG. 1.

 Thus, the need remains for a method for terminating tunnels initiated on multiple tunnel initiators on a common tunnel endpoint.

SUMMARY OF THE INVENTION

In accordance with preferred embodiments of the present invention, some of the problems associated with routing multiple communication links to the same endpoint are overcome.

5 An embodiment of a method, according to the present invention, for determining an endpoint in a communications system includes receiving a connection request from a client and, responsive to the connection request, querying a database for a database entry matching the client using predetermined identifying information, where the matching database entry will include an identifier for an endpoint. Responsive to receiving a database reply including the
10 identifier for the endpoint, the method then sets forth establishing a connection for the client to the endpoint identified in the database reply. Alternatively, responsive to not receiving a database reply, the method calls for establishing a connection for the client to a locally determined endpoint, and updating the database to include a database entry that includes the predetermined identifying information for the client and an identifier for the locally determined
15 endpoint.

The foregoing and other features and advantages of a preferred embodiment of the present invention will be more readily apparent from the following detailed description, which proceeds with references to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is described in the context of an embodiment of the invention with reference to the following drawings, wherein:

FIG. 1 is a functional block diagram illustrating a network architecture having two
5 prearranged tunnel connections that terminate on different endpoint devices;

FIG. 2 is a functional block diagram illustrating a network architecture according to an embodiment of the present invention having two prearranged tunnel connections that terminate on a single endpoint device, where a database device is directly coupled to the network;

FIG. 3 is a functional block diagram illustrating another network architecture according
10 to an embodiment of the present invention, where the database is locally connected to a tunnel initiator;

FIG. 4 is a functional block diagram illustrating yet another network architecture according to an embodiment of the present invention, where the database is locally shared to a cluster of tunnel initiators;

FIG. 5 is a message sequence scenario illustrating an example of message traffic
15 according to an embodiment of the present invention when a client call is received for which there is no database entry;

FIG. 6 is a message sequence scenario illustrating an example of message traffic according to an embodiment of the present invention when a client call is received and there is a
20 matching database entry;

FIG. 7 is a message sequence scenario illustrating an example of message traffic according to an embodiment of the present invention based upon multicasting when a client call is received and there is a matching database entry;

FIG. 8 is a message sequence scenario illustrating an example of message traffic according to an embodiment of the present invention based upon repeated multicasting when a client call is received and there is no matching database entry;

FIG. 9 is a message sequence scenario illustrating an example of message traffic according to an embodiment of the present invention when a call is disconnected;

FIG. 10 is a functional block diagram illustrating a mobile network architecture according to an embodiment of the present invention; and

FIG. 11 is a message sequence scenario illustrating an example of message traffic according to an embodiment of the present invention based upon multicasting when a client call from a mobile client is received and there is a matching database entry.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

The present invention is directed toward a method for terminating multiple tunnel connections on a common tunnel endpoint.

FIG. 2 is a block diagram illustrating a network architecture 100 suitable for use with the present invention. Architecture 100 includes a database 112 linked to IP network 70 for storing tunnel endpoint information for tunnel connections established for remote clients. Tunnel Initiators 130 and 140 are adapted to follow a protocol according to the present invention that requires the TIs to consult database 110 before establishing a call to a tunnel endpoint.

One example of a tunnel initiator device is a network access server, such as that described in the patent to Dale M. Walsh et al., U.S. No. 5,525,595, which is fully incorporated by reference herein and describes an integrated network access server suitable for use in the present invention. Such a device has been commercialized widely by 3Com Corporation (previously
5 U.S. Robotics Corp.) under the trade designation Total Control TM Enterprise Network Hub. Network access servers similar in functionality, architecture and design are available from other companies, including Ascend Communications, Livingston Enterprises, Multitech, and others. The invention is suitable for implementation in network access servers from the above companies, and other similar devices.

10 According to the protocol of the present invention, a TI will send an IP multicast-based query to database 110 that includes an <EDO, Username> pair. An Endpoint Discriminator (EDO) is a unique identifier for a device, such as a central processor unit identifier (CPUid) or a Medium Access Control (MAC) address. The Username is a value registered with a target Remote Access Server (RAS).

15 The IP multicast message will have a predetermined message type that uniquely identifies it as a database query in accordance with the present invention. IP multicasting is the transmission of an IP datagram to a "host group", a set of zero or more hosts identified by a single IP destination address. A multicast datagram is delivered to all members of its destination host group with the same "best-efforts" reliability as regular unicast IP datagrams, i.e., the
20 datagram is not guaranteed to arrive intact at all members of the destination group or in the same order relative to other datagrams. The membership of a host group is dynamic; that is, hosts may join and leave groups at any time. There is no restriction on the location or number of members

in a host group. A host may be a member of more than one group at a time and a host need not be a member of a group to send datagrams to it.

A host group may be permanent or transient. A permanent group has a well-known, administratively assigned IP address. It is the address, not the membership of the group, that is permanent; at any time a permanent group may have any number of members, even zero. Those IP multicast addresses that are not reserved for permanent groups are available for dynamic assignment to transient groups which exist only as long as they have members. See RFC 1112 and RFC 2236 for further information regarding IP multicasting.

Database 112 will contain tuples having a format <EDO, Username, EP>, where EP is the endpoint address for a tunnel endpoint corresponding to the EDO and Username values. Responsive to the query, database 110 will return a message indicating failure, i.e. no tuple was found matching the EDO and Username values, or success. A success message will include the value of EP from the matching tuple in database 110.

Note that while database 110 is shown as an entity connected to network 70, the database may reside elsewhere relative to the tunnel initiator. For instance, database 110 can also reside locally on a tunnel initiator, as reflected in another embodiment of a network architecture 200 illustrated in FIG. 3. Or, the tunnel initiator 140A may be part of a cluster of tunnel initiators 140A and 140B connected via a local network 312 where the database 110 is coupled to the tunnel initiators through the local network 312, as reflected in yet another embodiment of a network architecture 300 illustrated in FIG. 4.

The use of IP multicasting permits the location of database 110 to be transparent to the protocol according to the present invention since the multicast message is universally broadcast

and will therefore be received by the database regardless of where it is located. Also, multicasting allows database 110 to be distributed across multiple devices, since multicasting provides for the handling of multiple replies to a multicast message by discarding extraneous replies.

5 FIG. 5 illustrates a first scenario 330 that illustrates a sequence of messages according to the present invention that occur as a result of a request for the establishment of a first link of a tunnel, such as client 20 might send to tunnel initiator 30 over link 22. In scenario 330, each vertical line represents one or more entities, such as those that appear in FIGS. 2, 3 and 4. For example, CLIENT corresponds to client 20, TUNNEL INITIATOR corresponds to tunnel initiator 130, 140, 140A or 140B, ENDPOINT DATABASE corresponds to endpoint database 110, TUNNEL ENDPOINT corresponds to tunnel endpoint 50 or 60, and SERVER corresponds to server 80. One of ordinary skill in the art will readily recognize that the present invention may be applied to other configurations without departing from the teachings of the present invention.

10 In FIG. 5, a multi-link point-to-point (MLPPP) call 332 from a client to a tunnel initiator causes the tunnel initiator to generate a multicast query 334 containing an EDO and
15 USERNAME that identifies the client. The endpoint database receives the query and searches the database for a tuple having key attribute values that match the EDO and USERNAME. In scenario 330, this is the first link for the client, so no tuple exists and the endpoint database returns a FAILURE message 336. Alternatively, the FAILURE message 336 may be sent when
20 the endpoint database finds that an EP ADDR field for the matching tuple is a null value indicating that no connections exist for the client.

In a multicast based alternative, the FAILURE message 336 is not sent and the endpoint database discards the query message. In the multicast alternative, the tunnel initiator retransmits the query a predetermined number of times, or times-out, and, based on the absence of a response, determines that no matching tuple exists. A multicast-based embodiment of the present invention is described in greater detail below with respect to FIG. 8.

Once the tunnel initiator determines that no other link exists for the client, it sends a tunnel set-up message 338 to a locally determined tunnel endpoint. Once the tunnel is established, the tunnel initiator multicast a database update message 340 containing the EDO and USERNAME values for the client and the EP ADDR value for the tunnel endpoint. The endpoint database receives the update message 340 and stores the information in the matching tuple in the database. A first tunnel link is now in place for data transfer 342 from the client to the server.

FIG. 6 illustrates a scenario 350 where the same client requests a second tunnel link via MLPPP call 352. This time, database query 354 results in the endpoint database finding the matching tuple that contains the EP ADDR of the tunnel endpoint for the first link. The endpoint database returns the EP ADDR in a success message 356 to the tunnel initiator. The tunnel initiator then uses the EP ADDR from message 356 to send tunnel set-up message 358 to the tunnel endpoint corresponding to the value of EP ADDR. This establishes a second link of a MLPPP connection between the tunnel initiator and the tunnel endpoint and the multi-link connection is now in place for data transfer 360 between the client and server.

Alternatively, the tunnel initiator according to the present invention can be configured to check a local database for the matching tuple for the client EDO and USERNAME values before

sending a multicast query. Messaging scenario 370 of FIG. 7 illustrates an example of the message flow for such a tunnel initiator.

In response to a MLPPP call 372 from the client, the tunnel initiator sends a local query 374 that contains the EDO and USERNAME for the client. The local database can be resident upon the tunnel initiator, on another tunnel initiator in the same cluster, or otherwise locally accessible by the tunnel initiator receiving the MLPPP call 372. If the local database does not find a matching tuple, then it returns a failure message 376 to the tunnel initiator which then sends out the multicast query 380 to obtain the EP ADDR value from an endpoint database residing elsewhere in the network.

Once the tunnel for the connection is set-up via message 382, a multicast update 384 updates the non-resident endpoint database. However, the local database can also be updated so that another call set-up request from the client that is received by the tunnel initiator will result in the local query being successful.

FIG. 8 illustrates an embodiment of the protocol according to the present invention that is based upon UDP multicasting. In the FAILURE response based approach of FIG. 7, problems can arise if requests for two links from the same client are pending at the same time. Each request can receive a FAILURE message and each tunnel initiator will locally select a tunnel endpoint, which may well result in two separate tunnel endpoints being selected.

In a multicast-based approach, each tunnel initiator retransmits a multicast query a predetermined number of times without receiving a response before concluding that no tunnel endpoint is determined for the client. If one tunnel initiator reaches its maximum number of retransmissions without a response before the other, then it will locally select a tunnel endpoint

and update the tunnel endpoint table with its selection. When the other tunnel initiator retransmits its query, it will receive a reply based upon the information in the first tunnel initiator's update.

To be compatible with UDP multicasting, each server that hosts an endpoint database or a portion thereof must be configured to silently discard a multicast query for which it has no matching tuple. In scenario 390 of FIG. 8, the tunnel initiator, in response to MLPPP call 392, generates a predetermined number of multicast queries 394A, 394B and 394C. The tunnel initiator waits a predetermined time period for a response after each query attempt.

When no response is received after the last query 394C is sent, then the tunnel initiator determines that no tuple exists in the endpoint database, or databases, and proceeds to establish the tunnel connection using a locally determined tunnel endpoint. The tunnel initiator then sends out a multicast update 398 that updates the endpoint database or databases. With multiple databases, a subsequent multicast query with the clients EDO and USERNAME values will result in multiple responses. The tunnel initiator is configured to accommodate multiple responses by discarding the extraneous responses.

The protocol according to the present invention can also be applied to mobile networks, as mentioned above. FIG. 10 shows an embodiment of a mobile architecture 410 that illustrates the application of the present invention to mobile communications. Mobile client 420 has a wireless connection 422 to tunnel initiator 430, which is, for example, a cell site. Tunnel initiator 432 has established a connection 432 to tunnel endpoint 450 in order to communicate with server 80.

However, when mobile user 420 leaves the service area for tunnel initiator 430, another link 424 must be established with the tunnel initiator 440 for the service area that the mobile user is entering. It is advantageous to terminate the tunnel connection for link 424 to the same tunnel endpoint 450 as terminates tunnel connection 456 for link 422.

5 In mobile IP, mobile user 420 typically senses that it has lost contact with tunnel initiator 430 when it times out waiting for an advertisement message from tunnel initiator 430. Mobile user 420 will then look for an advertisement message 444 from a new foreign agent, tunnel initiator 440. When mobile user 420 receives advertisement message 444, it sends a registration message 426 that initiates set-up of link 424 and set-up of tunnel connection 466. See RFC 2002
10 for further details regarding mobile IP.

When tunnel initiator 440 receives the registration message 426, it follows the protocol according to the present invention in establishing tunnel connection 466 to tunnel endpoint 450. FIG. 11 illustrates a scenario 490 for the protocol as applied to the mobile network architecture 410 of FIG. 10. The registration call, 426 in FIG. 10 and 492 in FIG. 11, prompts tunnel initiator
15 440 to broadcast database query 494 containing the EDO and USERNAME of mobile user 422. Endpoint database 110 receives the query and searches for a matching tuple. A matching tuple with the EDO, USERNAME and EP ADDR for tunnel endpoint 450 will exist in the database because tunnel initiator 430 will have sent an update to database 110 after having established tunnel connection 456 to endpoint 450 for link 422. Endpoint database 110 will return a success
20 message 496 containing the EP ADDR for tunnel endpoint 450 and tunnel initiator 440 will establish tunnel connection 466 with endpoint 450.

Tear-down of connections and clean up of entries in endpoint database 110 can occur in a variety of ways. Database entries may include a timestamp that provides for entries to be removed from the database after a pre-determined time period. Database entries may also be removed responsive to a tear-down message from the tunnel initiators or tunnel endpoint or, in the case of mobile IP, through a de-registration message. In the latter case, the message will include a unique tear-down message type along with an identifier for the database entry, such as the EDO/USERNAME combination or a mobile identification number (MIN).

The protocol according to the present invention supports the deterministic selection of an endpoint for connections having multiple origination points. Although the present invention is described in the context of an L2TP tunnel and a mobile connection, the present invention is applicable to any communications link where it is desirable to terminate connections from multiple origins to the same endpoint.

It should be understood that the programs, processes, methods, systems and apparatus described herein are not related or limited to any particular type of computer apparatus (hardware or software), unless indicated otherwise. Various types of general purpose or specialized computer apparatus may be used along with the present invention or perform operations in accordance with the teachings described herein.

In view of the wide variety of embodiments to which the principles of the invention can be applied, it should be understood that the illustrated embodiments are exemplary only, and should not be taken as limiting the scope of the present invention. For example, the messages of the message flow scenarios may be taken in sequences other than those described, and more or

fewer elements or components may be used in the block diagrams. In addition, the present invention can be practiced with software, hardware, or a combination thereof.

The claims should not be read as limited to the described order or elements unless stated to that effect. Therefore, all embodiments that come within the scope and spirit of the following

5 claims and equivalents thereto are claimed as the invention.

WE CLAIM:

1. A method for determining an endpoint in a communications system, the method comprising the steps of:

receiving a connection request from a client;

responsive to the connection request, querying a database for a database entry matching the client using predetermined identifying information, where the matching database entry will include an identifier for an endpoint;

responsive to receiving a database reply including the identifier for the endpoint, establishing a connection for the client to the endpoint identified in the database reply; and

responsive to not receiving a database reply:

establishing a connection for the client to a locally determined endpoint, and

updating the database to include a database entry that includes the predetermined identifying information for the client and an identifier for the locally determined endpoint.

2. The method of claim 1, where the step of querying a database for a database entry matching the client includes querying a local database for the matching entry and, if none is found, querying a remote database.

3. The method of claim 2, where the step of querying a remote database further includes multicasting a message having a predetermined message identifier and the predetermined identifying information for the client.

4. The method of claim 3, where:

the step of multicasting a message having a predetermined message identifier further includes multicasting the message having the predetermined message identifier a predetermined
5 number of times in the absence of a database reply; and

the step, responsive to not receiving a database reply, of establishing a connection for the client to a locally determined endpoint further includes waiting a predetermined time-out period after a last multicast of the predetermined number of times for multicasting before establishing the connection for the client to the locally determined endpoint.

10
5. The method of claim 4, where the step of updating the database to include a database entry that includes the predetermined identifying information for the client and an identifier for the locally determined endpoint further comprises multicasting a message having another predetermined message identifier and including the predetermined identifying
15 information for the client and the identifier for the locally determined endpoint.

6. The method of claim 5, where the predetermined identifying information includes an EDO and a Username for the client.

20 7. The method of claim 6, where the identifier for the endpoint is a network address.

8. The method of claim 1, further including the steps of:

receiving another connection request from the client;

responsive to the another connection request and using the predetermined identifying information for the client, querying the database for the database entry that includes the predetermined identifying information for the client and the identifier for the locally determined endpoint; and

responsive to the database reply including the identifier for the locally determined endpoint, establishing another connection for the client to the locally determined endpoint.

9. A computer readable medium having stored therein instructions for causing a central processing unit to execute the method of claim 1.

10. A network communication system, the system comprising:

a database device configured to store a data entry, where the data entry is keyed by predetermined client identifying information and includes an endpoint identifier field, the database device being further configured to receive a database query that includes a client identifying information value, search for a matching data entry that matches the client identifying information value and, if the matching data entry is found, send a database reply that includes the value of the endpoint identifier field of the matching data entry;

a network device for receiving a call request from a client and, responsive thereto, generate a database query having the client identifying information value for the client from which the call request is received, and where the network device is further configured, when a database reply corresponding to the database query for the calling client is received, to establish a connection to a network device corresponding to the endpoint identifier value included in the database reply and, when no database reply corresponding to the database query for the calling client is received, the network device is configured to locally select a locally determined endpoint value and establish a connection for the client to a network device corresponding to the locally determined endpoint value.

11. The network communication system of claim 10, wherein:

the network device is further configured to generate a database update when it establishes the connection for the client to the network device corresponding to the locally determined endpoint value, where the database update includes the client identifying information value for the client from which the call request is received and the locally determined endpoint value; and

the database device is further configured, responsive to the database update, to create a database entry having the client identifying information value for the client from which the call request is received and the locally determined endpoint value.

5 12. The network communication system of claim 10, where the network device is further configured to generate the database query having the client identifying information value for the client from which the call request is received by sending a multicast message having a predetermined message type and the client identifying information value for the client from which the call request is received and where the database device is configured to search for the
10 matching database entry responsive to receiving the multicast message having the predetermined message type.

13. The network communication system of claim 12, where the network device is further configured to determine that no database reply has been received by sending the multicast
15 message having the predetermined message type a predetermined number of times without receiving the database reply within a predetermined time period.

14. The network communication system of claim 11, where:

the network device is further configured to generate the database update by sending a
20 multicast message having a predetermined message type along with the client identifying information value for the client from which the call request is received and the locally determined endpoint value; and

the database device is configured to create the database entry having the client identifying information value for the client from which the call request is received and the locally determined endpoint value responsive to receiving the multicast message having the predetermined message type.

5

15. The network communication system of claim 11, where the client identifying information includes an EDO and a Username for the client and the locally determined endpoint value is a network address configured to route the connection through a network to the network device corresponding to the locally determined endpoint value.

10

16. The network communication system of claim 11, the system further including another network device for receiving another call request from the client and, responsive thereto, generate another database query having the client identifying information value for the client from which the another call request is received, and where the another network device is further configured, when a database reply corresponding to the another database query for the calling client is received, to establish another connection to the network device corresponding to the endpoint identifier value included in the database reply corresponding to the another database query.

15

20

17. A network system for transmitting and receiving packets across a network, the network device comprising:

a database device coupled to the network and configured to receive a first predetermined type of message having a client identifier field and, responsive thereto, search for a database entry having a key field with a value matching a value of the client identifier field, where the database entry includes an endpoint address field and, when a corresponding database entry is found, generate a second predetermined type of message having the client identifier field that includes the value of the key field of the database entry and an endpoint field that includes a value of the endpoint address field of the database entry, and where the database device is configured to receive a third predetermined type of message having the client identifier field and the endpoint field and, responsive thereto, store a database entry corresponding to a value in the client identifier field of the third predetermined type of message and having a value of the endpoint field of the third predetermined type of message in the endpoint address field of the database entry;

a first network device coupled to the network and configured to receive a first call request from a client device and, responsive thereto, locally select a second network device coupled to the network and establish a first connection for the client device from the first network device to the second network device, where the first network device is also configured to generate the third predetermined type of message having a value corresponding to the client device in the client identifier field and an address value corresponding to the second network device in the endpoint field; and

a third network device coupled to the network and configured to receive a second call request from the client device and, responsive thereto, generate the first predetermined type of message having the value corresponding to the client device in the client identifier field, and where the third network device is also configured to wait for a predetermined time period to receive the second predetermined type of message having a value of the client identifier field corresponding to the client device and, using the value of the endpoint field of the second predetermined type of message, establish a second connection to the second network device.

18. The network system of claim 17, where the first predetermined type of message is a database query message, the second predetermined type of message is a database reply message, and the third predetermined type of message is a database update message.

19. The network system of claim 17, where:
the first network device is coupled to the third network device and the first network device is further configured to locally store a local database entry having a key field containing the value corresponding to the client device and an endpoint field containing the address value of the second network device; and

the third network device is further configured to query the first network device for the local database entry in response to receiving the second call request from the client device.

ABSTRACT OF THE DISCLOSURE

A method is shown for deterministically selecting an endpoint for multiple links through a network for a client device. When a first network device establishes a connection for a client device to an endpoint device, it causes a database entry to be created that is keyed by a client identifier value for the client device and includes an endpoint address field with the network address of the endpoint device. When a second network device receives a call request from the client device, it queries the database using the client identifier value for the client device and waits a predetermined period for a database reply. A database device storing the database entry receives the query from the second network device, searches for the database entry matching the client identifier value, and sends a database reply that includes the client identifier value and the network address of the endpoint device from the database entry. The second network device receives the database reply and uses the network address from the database reply to establish a second connection to the endpoint device for the client.

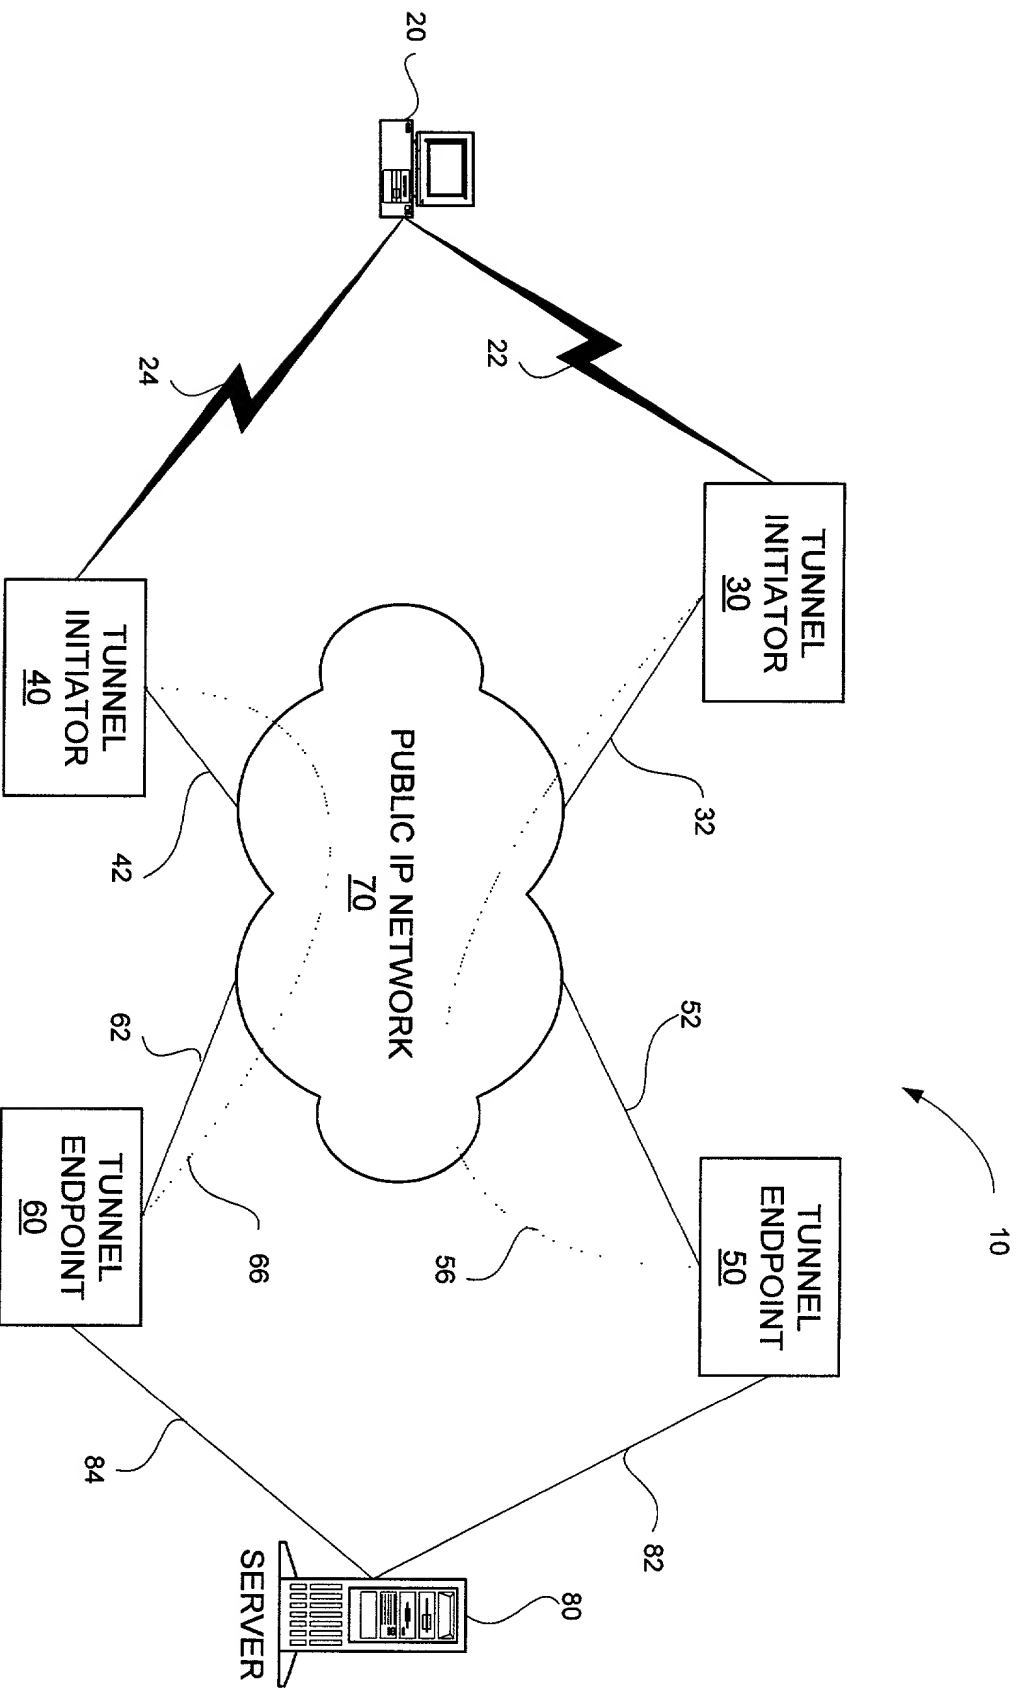


FIGURE 1

00430227-1000000

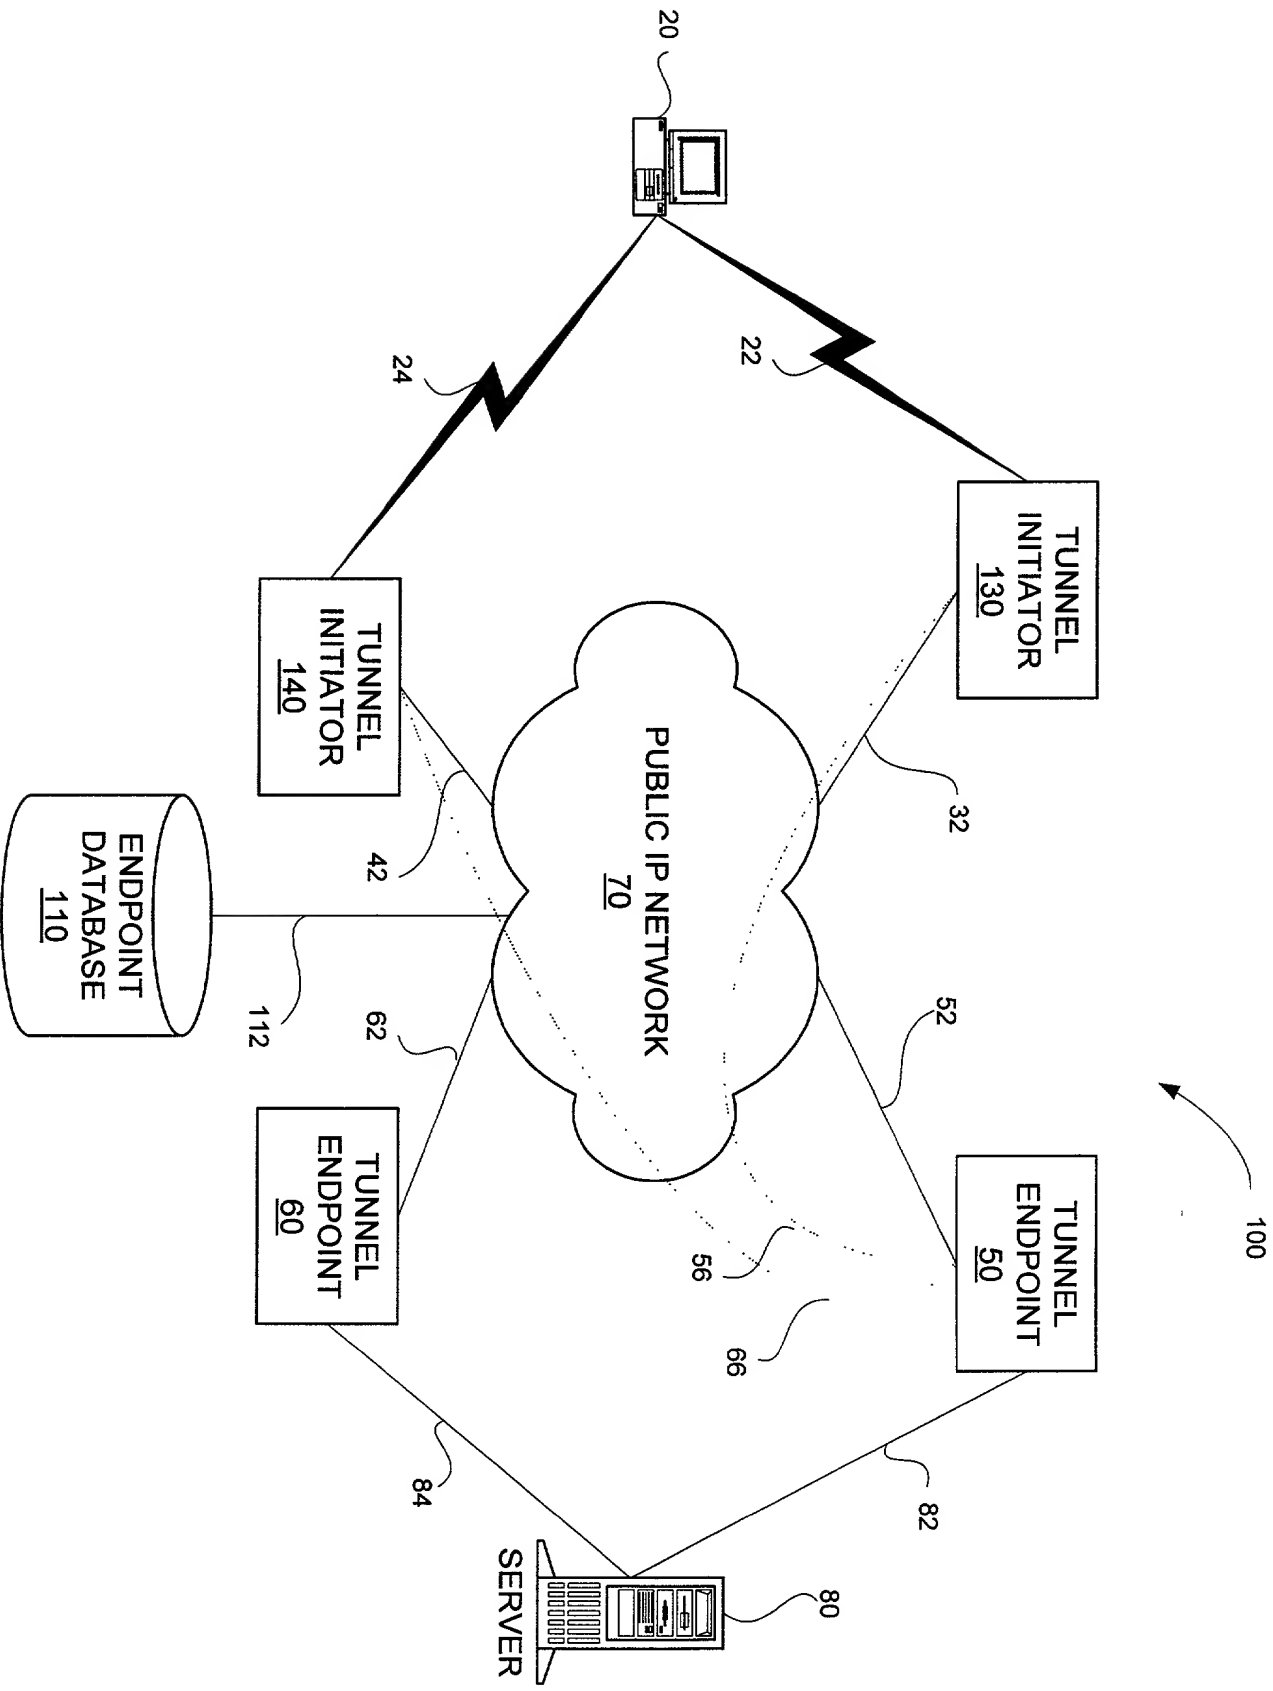


FIGURE 2

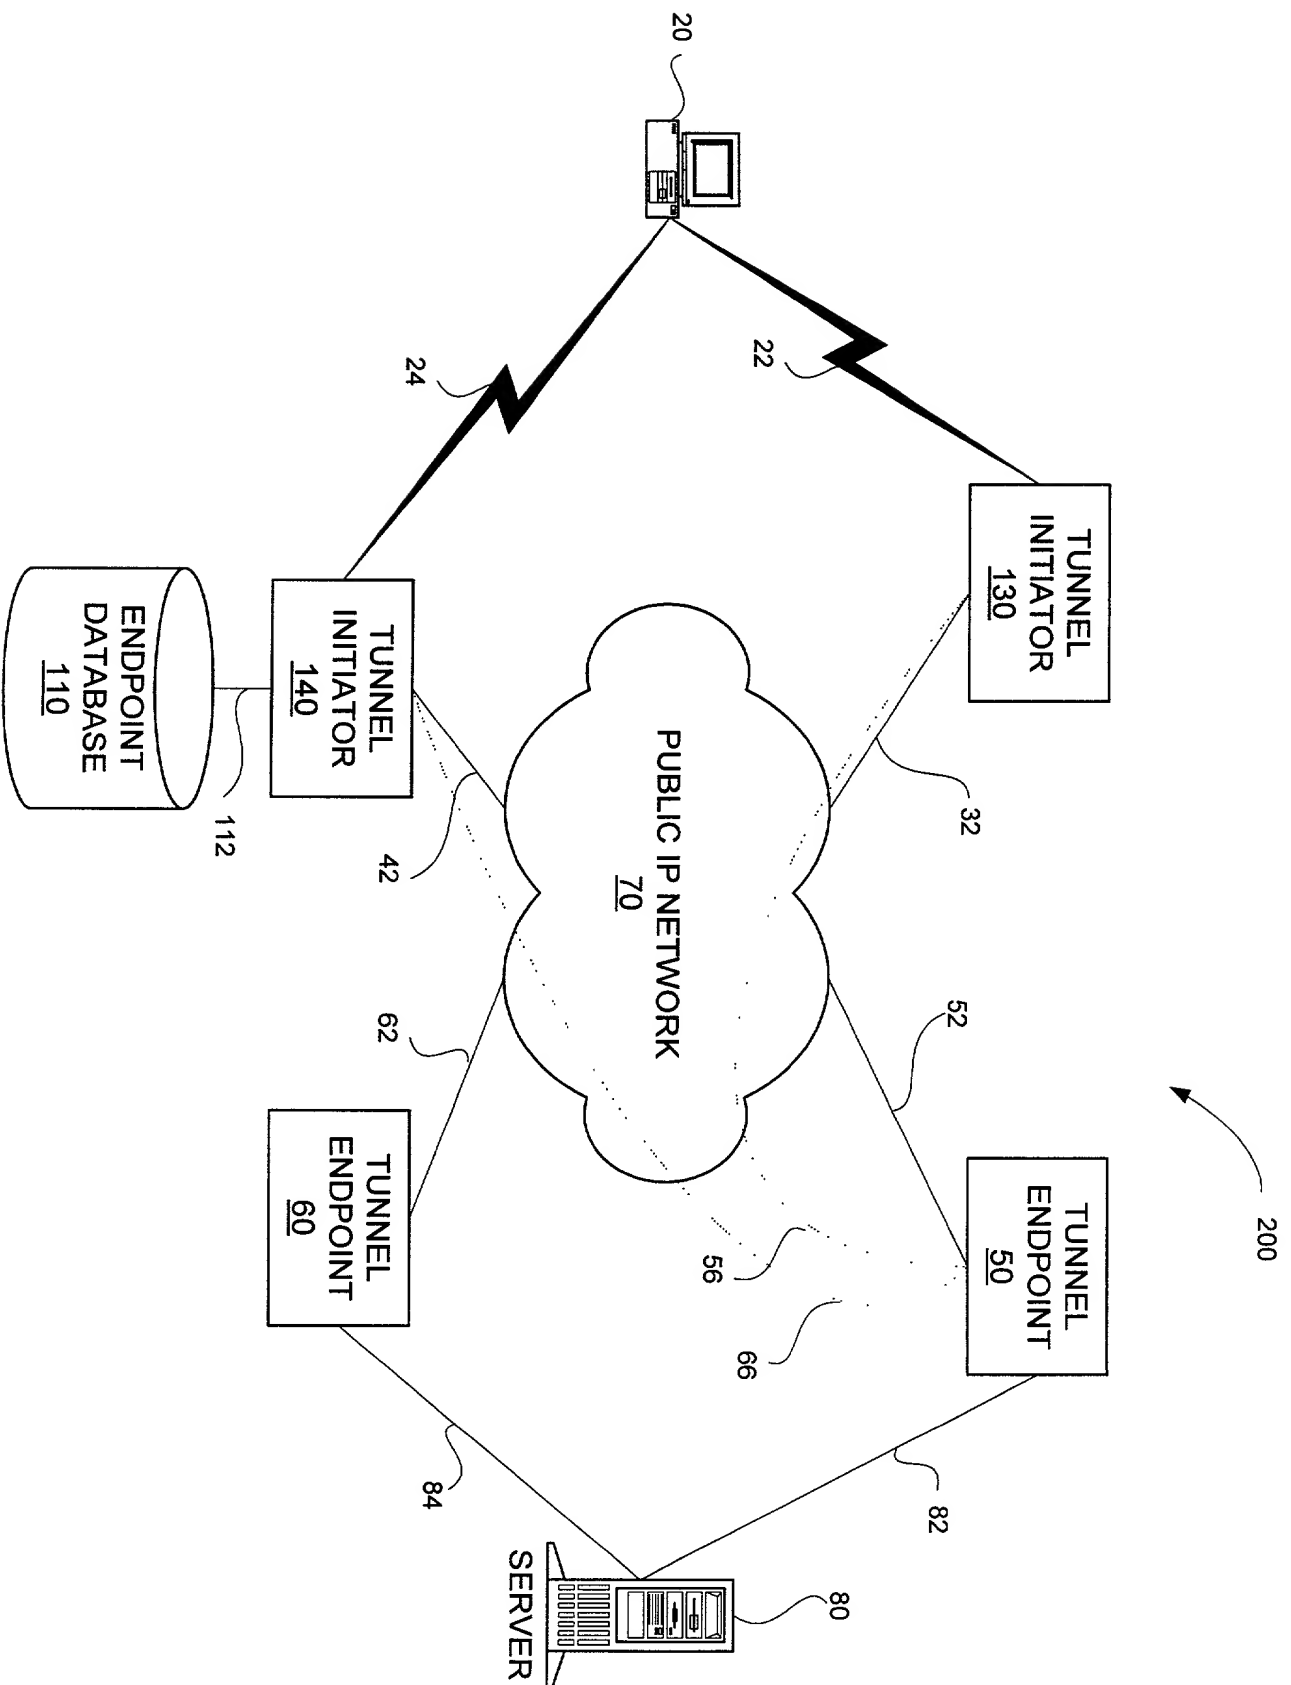


FIGURE 3

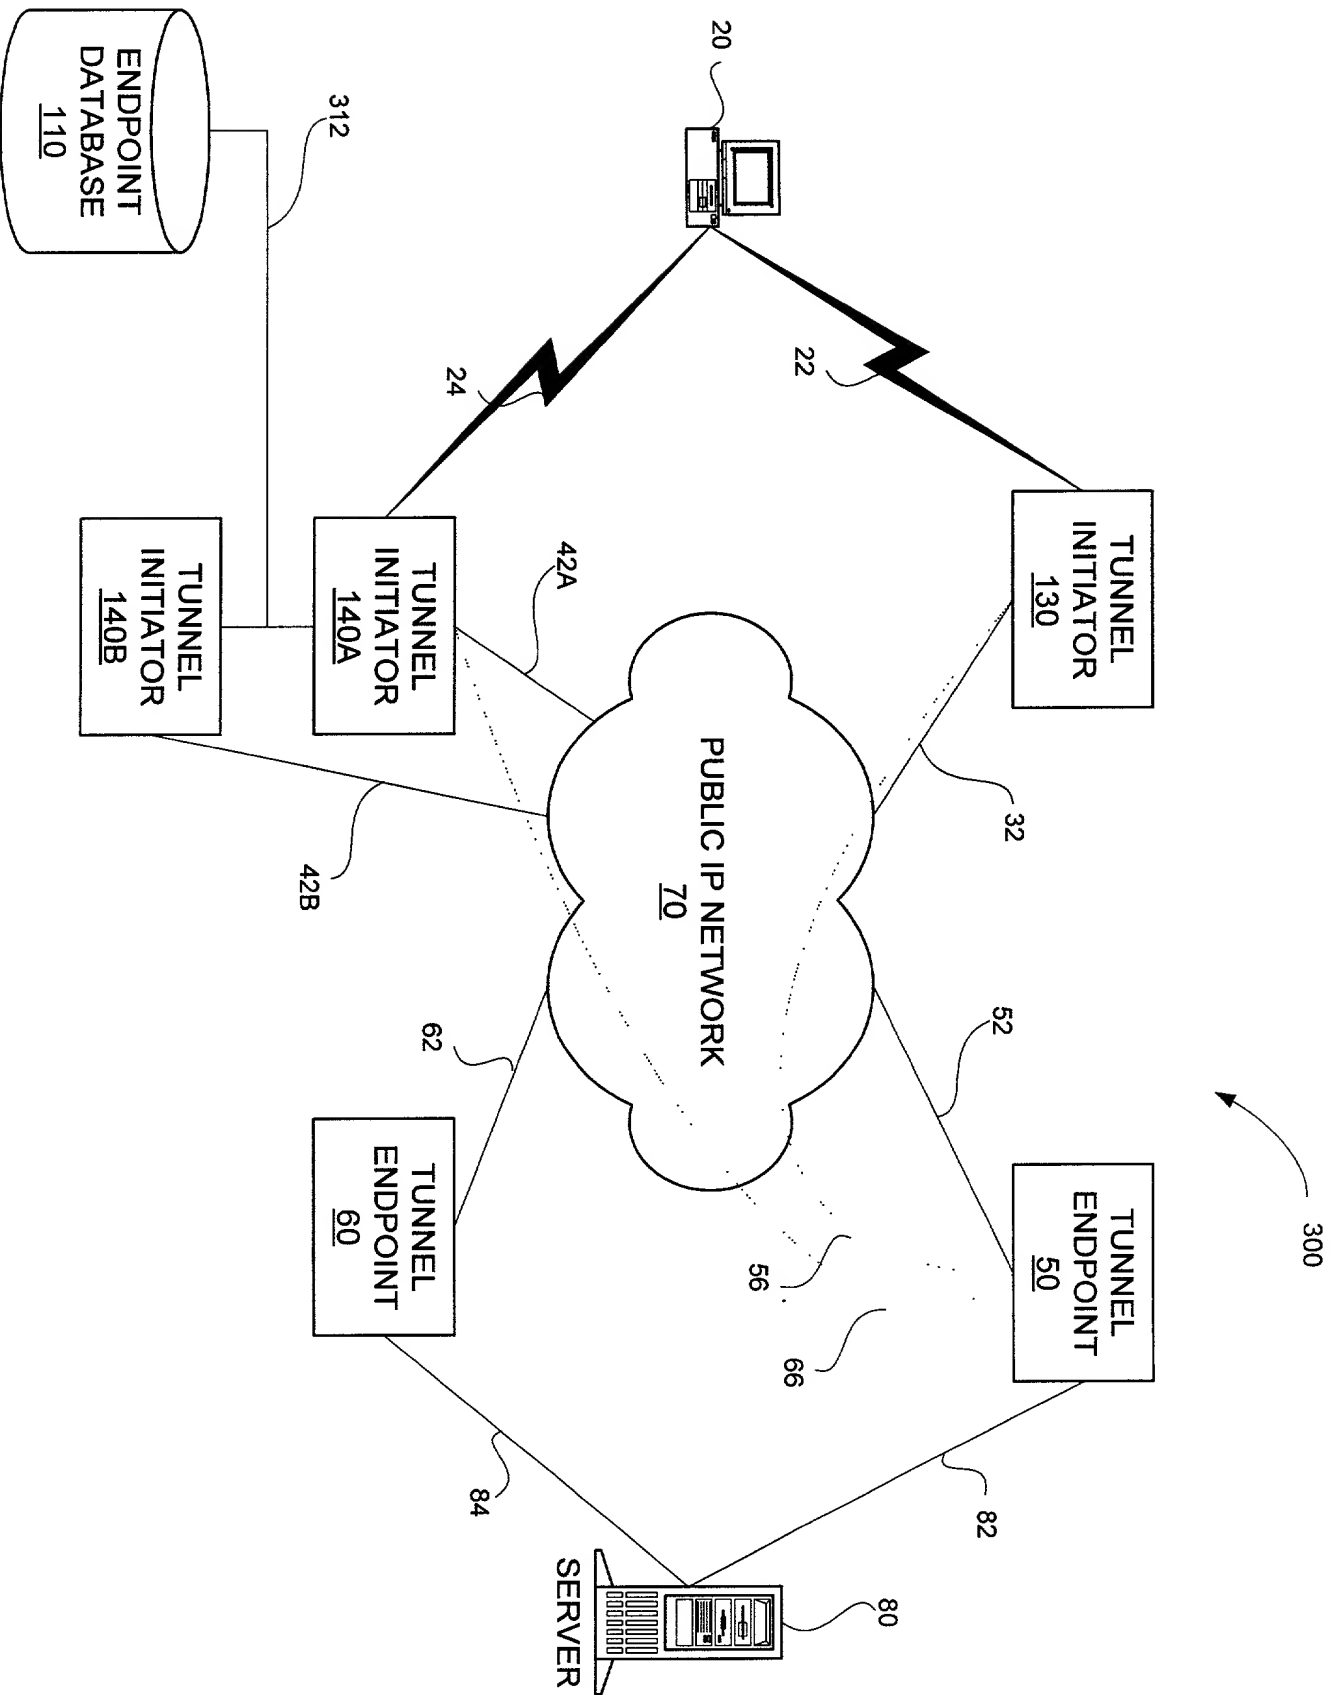


FIGURE 4

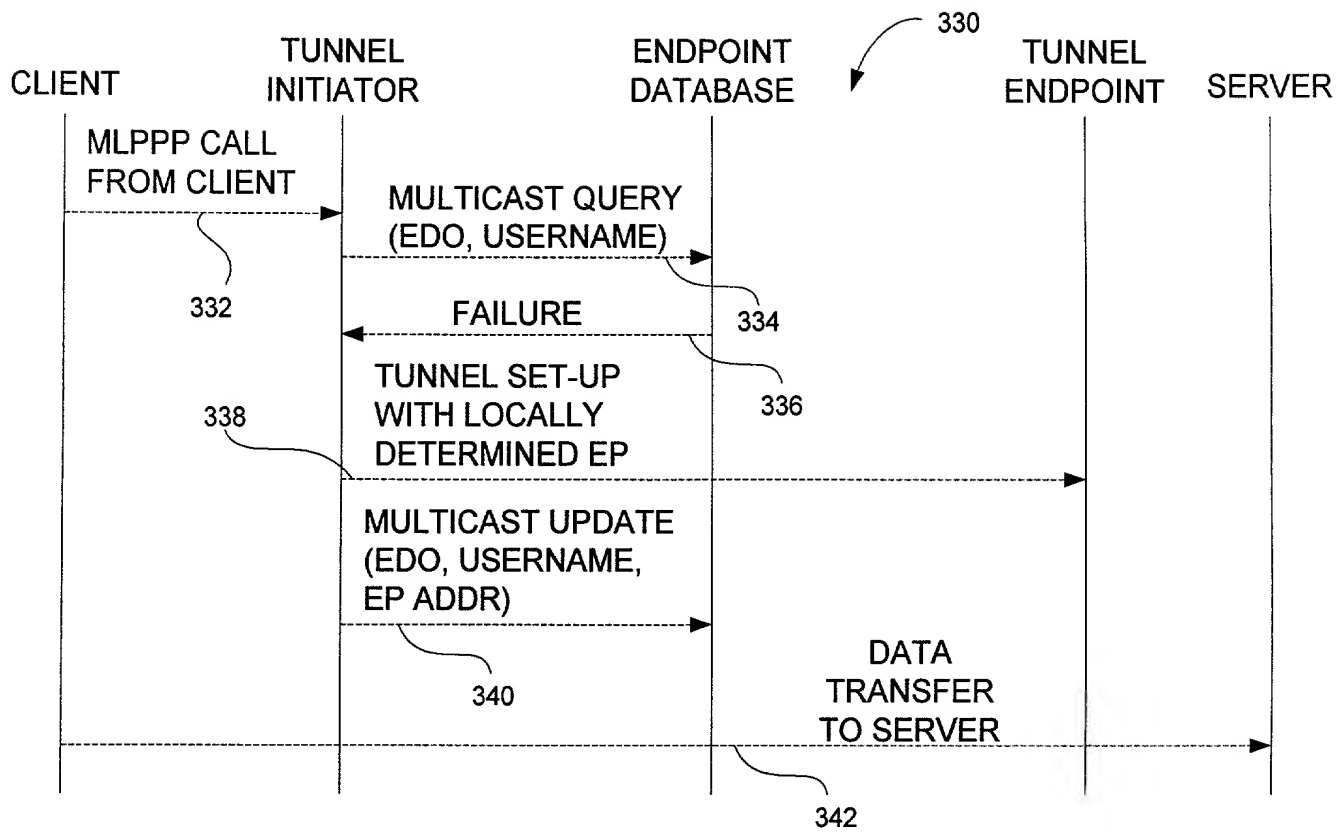


FIGURE 5

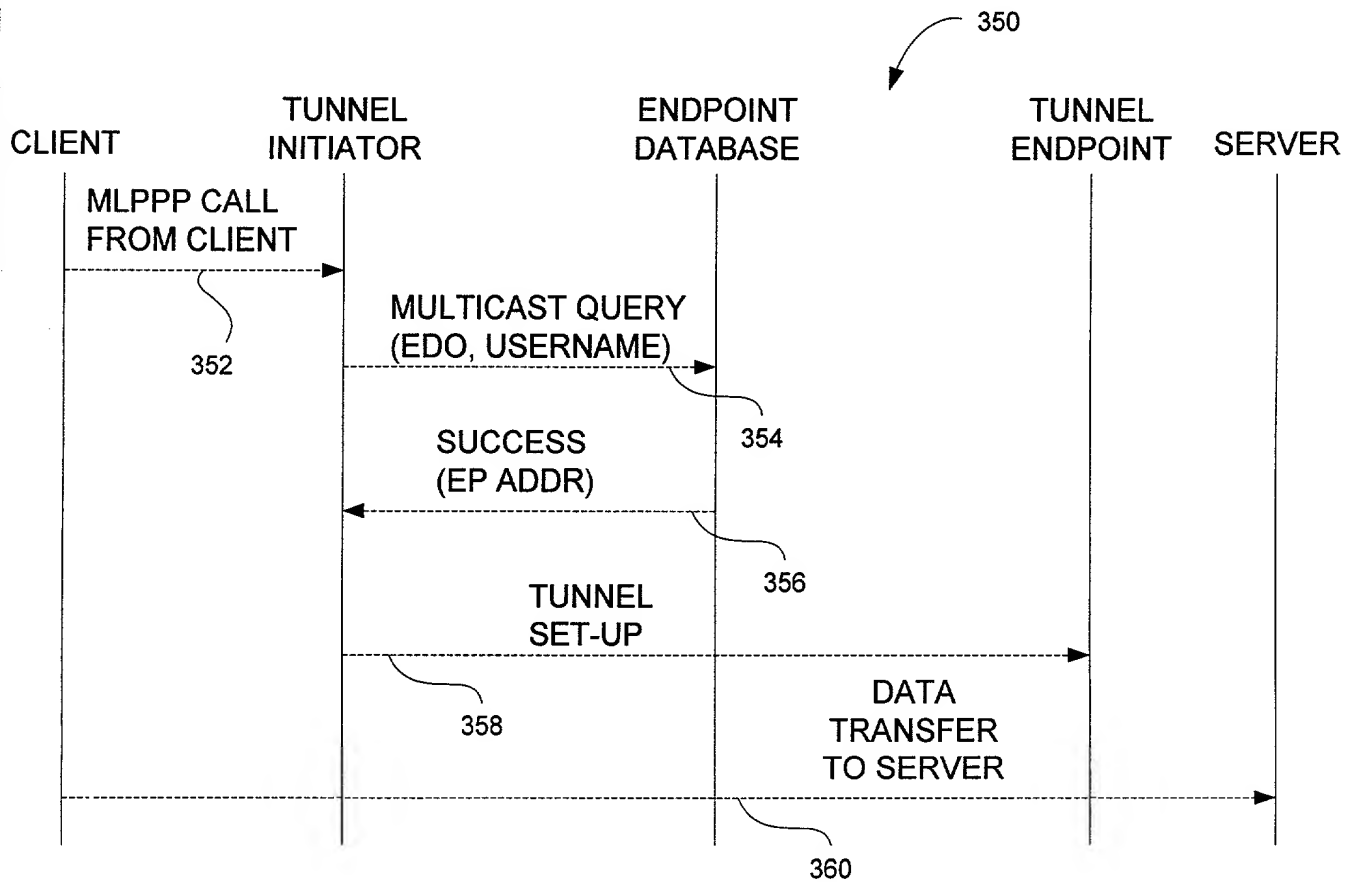


FIGURE 6

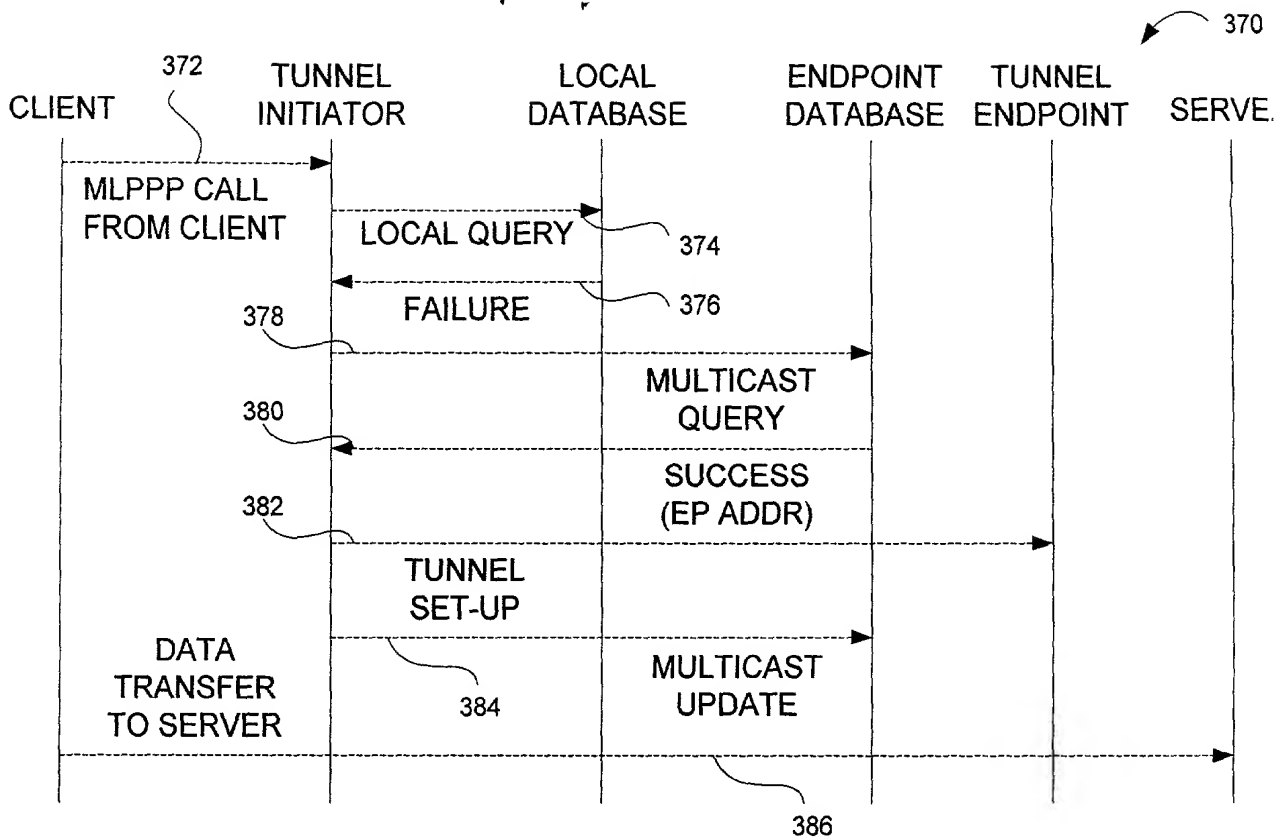


FIGURE 7

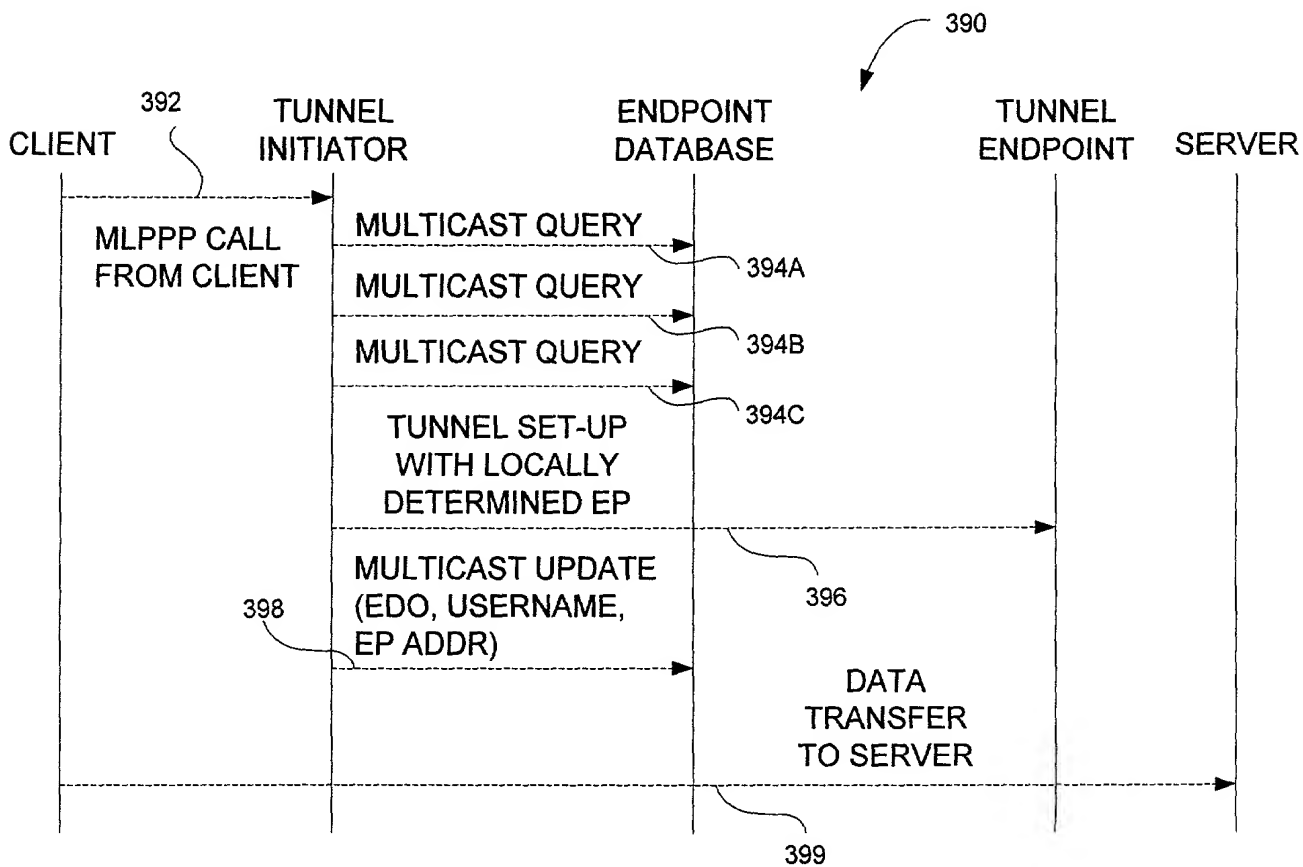


FIGURE 8

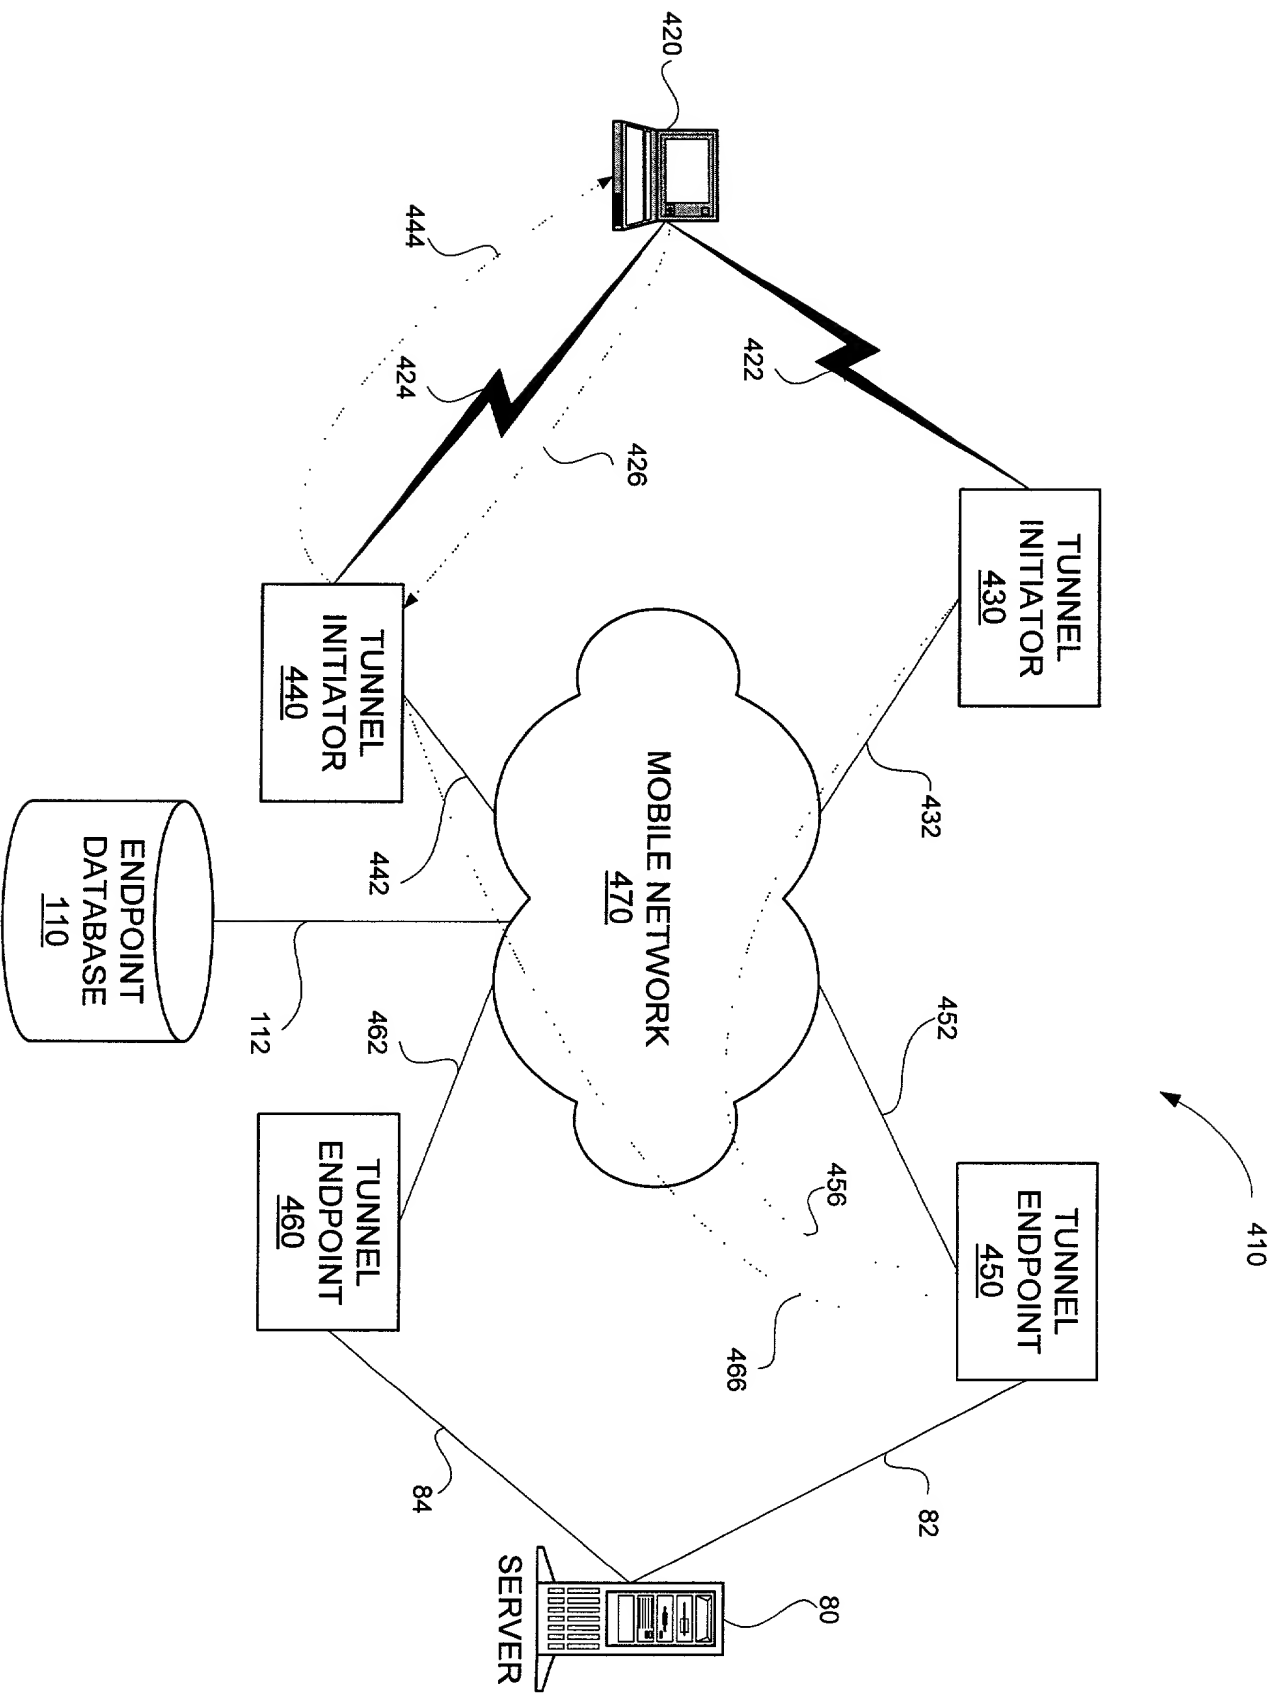


FIGURE 10

